

Sequential Logistic Principal Component Analysis (SLPCA): Dimensional Reduction in Streaming Multivariate Binary-State System

Zhaoyi Kang, Costas J. Spanos

Dept. of Electrical Engineering & Computer Sciences, UC Berkeley, Berkeley, CA 94709
{kangzy, spanos}@berkeley.edu,

Abstract—Sequential or online dimensional reduction is of interests due to the explosion of streaming data based applications and the requirement of adaptive statistical modeling, in many emerging fields, such as the modeling of energy end-use profile. Principal Component Analysis (PCA), is the classical way of dimensional reduction. However, traditional Singular Value Decomposition (SVD) based PCA fails to model data which largely deviates from Gaussian distribution. The Bregman Divergence was recently introduced to achieve a generalized PCA framework. If the random variable under dimensional reduction follows Bernoulli distribution, which occurs in many emerging fields, the generalized PCA is called Logistic PCA (LPCA) [1]. In this paper, we extend the batch LPCA to a sequential version (i.e. SLPCA), based on the sequential convex optimization theory. The convergence property of this algorithm is discussed compared to the batch version of LPCA (i.e. BLPCA), as well as its performance in reducing the dimension for multivariate binary-state systems. Its application in building energy end-use profile modeling is also investigated.

I. INTRODUCTION

Sequential data mining has received considerable attention recently as the development in wireless-sensor information technology facilitates the collection of huge amount of streaming data – This brings about several challenges on the efficiency in computation, storage and statistical learning [2]. Dimensional reduction in the streaming environment is one of the techniques that can help to overcome those issues [3].

Among the dimensional reduction techniques, Principal Component Analysis (PCA) is most widely-known. PCA finds the linear projection of the original data matrix which explains the largest portion of the variance. From the maximum likelihood perspective, PCA contains the assumption that the data follows Gaussian distribution. Naturally, this will fail to give reliable results when data largely deviates from Gaussian distribution [4]. Bregman Divergence is introduced to achieve a generalized PCA framework for a family of exponential distributed data (i.e. e PCA) [5]. As a generalization over the Frobenius norm, KL-divergence, Mahalanobis distance etc., Bregman Divergence is believed to better quantify the distance of variables coming from non-Gaussian distribution [6] [7]. In the case of Bernoulli random variables, which we are interested in, the generalized PCA can be viewed as Logistic PCA (LPCA).

In this work, we extend the LPCA to the sequential version, based on the sequential convex optimization theory [8] [9]. The convergence property of this algorithm is discussed with

respect to the batch optimization algorithm. An application in building energy end-use profile modeling is investigated as an experiment of this method, which demonstrates its capability in reducing dimension in multivariate binary-state systems.

This paper is organized as follows: In Section II, the background and the detail of the algorithm is given, including PCA, exponential family, the Bregman Divergence and eventually the sequential LPCA (i.e. SLPCA) which we propose. In Section III, the convergence property of the algorithm is discussed, followed by the simulation results as well as the application in energy end-use modeling in Section IV. In Section V, conclusion is drawn.

II. ALGORITHM FRAMEWORK

PCA as a dimensional reduction technique has been well studied, and our Sequential LPCA is essentially a generalized incremental version of the classical model.

A. Principal Component Analysis

PCA is a well-known technique for dimensional reduction for high dimension data. It is of special importance in high dimensional regression model, and in a variety of applications, ranging from face recognition to generalized machine learning [10] [2].

There are two perspectives of PCA [4]. The first is the matrix factorization perspective. For a matrix $\mathbf{X} \in \mathbb{R}^{N \times P}$, we find a lower rank matrix Θ to minimize the error:

$$\min_{\Theta} \|\mathbf{X} - \Theta\|_F^2 \quad (1)$$

in which $\|\cdot\|_F$ is the Frobenius norm. This problem can be solved by Singular Vector Decomposition (SVD).

Definition 1 (Singular Value Decomposition (SVD)). *for input data matrix $\mathbf{X} \in \mathbb{R}^{N \times P}$, Singular Value Decomposition decomposes the matrix to be:*

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_{i=1}^{\min(N,P)} \sigma_i \mathbf{u}_i \mathbf{v}_i^T \quad (2)$$

in which \mathbf{U} is an $N \times N$ matrix called column eigenvector matrix, with \mathbf{u}_i as i^{th} columns; $\mathbf{V} \in \mathbb{R}^{P \times P}$ called row eigenvector matrix, with \mathbf{v}_i as i^{th} columns; $\mathbf{\Sigma} \in \mathbb{R}^{N \times P}$ rectangular diagonal matrix, with i^{th} diagonal value as σ_i .

If σ_i 's are sorted and the largest is σ_1 , then $\Theta = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T$ is the solution to Equation (1) if $\text{rank}(\Theta) = 1$. If the $\text{rank}(\Theta) = r$, then we choose $\Theta = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T$.

However, there is another perspective of PCA that is less widely-known, which is called the probabilistic interpretation. Here, the columns of $\mathbf{X} \in \mathbb{R}^{N \times P}$ can be viewed as N samples drawn from a Gaussian distribution with dimension lower than P . This idea can be used in larger family of distributions, for example, the exponential family distributions.

B. Exponential Family

Definition 2 (Exponential Family). *In the exponential family of distributions the conditional probability of a value x given parameter value θ takes the following form [11]:*

$$\log P(x|\theta) = \log P_0(x) + x\theta - G(\theta) \quad (3)$$

In which, θ is the natural parameter of the distribution. $G(\theta)$ is a function that ensures that the sum (integral) of $P(x|\theta)$ over the domain of x is one. It is observed that $G(\theta) = \log \sum_x P_0(x) e^{x\theta}$.

Borrowing idea from the probabilistic view of PCA, if x is the original data, θ comes from a lower dimensional space.

C. Exponential Family PCA

Equation (1) becomes inappropriate when the data is not Gaussian, which happens a lot in real world. Instead of Frobenious norm, we need another way to quantify the distance between and its lower rank approximation. The Bregman Divergence is introduced to generalize the distance quantification [7] [6] [5].

Definition 3 (Bregman Divergence). *The Bregman divergence w.r.t. F is defined, for $p, q \in \mathbb{R}^d$, as:*

$$B_F(p||q) = F(p) - (F(q) + \nabla F(x)^T \cdot (p - q)) \quad (4)$$

For an exponential family distribution in (4), let $F(g(\theta)) + G(\theta) = g(\theta)\theta$, in which $g(x) = \nabla G(x)$. If \mathbf{P}, \mathbf{Q} are matrices, $B_F(\mathbf{P}||\mathbf{Q}) = \sum_{i,j} B_F(\mathbf{P}_{ij}||\mathbf{Q}_{ij})$.

Example 1. *In the case of Gaussian distribution, the Bregman Divergence equals to squared loss $B(x||g(\theta)) = \frac{1}{2}(x - \theta)^2$.*

Example 2. *In the case of Bernoulli distribution, Bregman Divergence is the logit function $B(x||g(\theta)) = \log(1 + \exp(-x^*\theta))$, in which x^* is a transformation of x as $x^* = 2x - 1 \in \{-1, 1\}$. In this case, Bregman Divergence is a convex function of θ , thus can be placed in an efficient optimization framework.*

Therefore, similar to Equation (1), we can construct an optimization problem based on the Bregman Divergence. For data matrix \mathbf{X} and $[\mathbf{X}]_{ij} = x_{ij}$, we use $g(\Theta)$ to approximate, i.e. $[\Theta]_{ij} = \theta_{ij}$:

$$\min_{\Theta} \sum_{i,j} B(x_{ij}||g(\theta_{ij})) \quad (5)$$

In this work, we mainly focus on Bernoulli random variables, in other words, the Logistic PCA (LPCA). Hence, the optimization problem is as in Example 2. The logit function in

some cases is not strictly convex, thus we need to regularize the θ variable by $L(\theta)$:

$$\min_{\Theta} \sum_{i,j} B(x_{ij}||g(\theta_{ij})) + L(\Theta) \quad (6)$$

If we want to optimize Equation (6) with the constraint of the rank of Θ , we can re-write Θ as $\mathbf{A}\mathbf{V}^T$ where $\mathbf{A} \in \mathbb{R}^{N \times r}$ and $\mathbf{V} \in \mathbb{R}^{P \times r}$ s.t. $\text{rank}(\mathbf{A}\mathbf{V}^T) = r$. Then, we will minimize over two matrices \mathbf{A} and \mathbf{V} .

$$\min_{\mathbf{A} \in \mathbb{R}^{N \times r}, \mathbf{V} \in \mathbb{R}^{P \times r}} \sum_{i,j} B(\mathbf{X}_{ij}||g(\mathbf{A}\mathbf{V}^T)) + \gamma \Gamma(\mathbf{A}) + \lambda R(\mathbf{V}) \quad (7)$$

where $\Gamma(\cdot)$ and $R(\cdot)$ are regularization functions. In our work, both functions are quadratic, $\Gamma(\cdot) = R(\cdot) = \frac{\|\cdot\|_F^2}{2}$.

D. Batch Logistic PCA (BLPCA)

For the optimization problem in Equation (7), we can solve it in an alternating minimization algorithm [5] [12]. If we define \mathbf{a}_t as the t^{th} row of \mathbf{A} , let:

$$h_t(\mathbf{a}_t, \mathbf{V}) = B(\mathbf{x}_t||g(\mathbf{a}_t \mathbf{V}^T)), t = 1, \dots, N \quad (8)$$

Then we can solve Equation (7) by iterating the following two steps, and we call this method Batch LPCA (BLPCA):

$$\mathbf{a}_t^* = \arg \min_{\mathbf{a} \in \mathbb{R}^{1 \times r}} h_t(\mathbf{a}, \mathbf{V}^*) + \gamma \frac{\|\mathbf{a}\|_F^2}{2}, \forall t \quad (9)$$

$$\mathbf{V}^* = \arg \min_{\mathbf{V} \in \mathbb{R}^{P \times r}} \sum_{t=1}^N h_t(\mathbf{a}_t^*, \mathbf{V}) + \lambda \frac{\|\mathbf{V}\|_F^2}{2} \quad (10)$$

E. Sequential Logistic PCA (SLPCA)

For a sequential version of BLPCA, \mathbf{V} is of fixed dimension when data is streaming in. However, the dimension of \mathbf{A} would change after every step. Similar to [13], at each time t , we solve a local sub-optimal for the t^{th} row of \mathbf{A} (i.e. $\tilde{\mathbf{a}}_t$) instead of a global one, and sequentially update \mathbf{V} with the $\tilde{\mathbf{a}}_t$'s (i.e. $\tilde{\mathbf{V}}^t$). At step t , this means that we solve for $\tilde{\mathbf{a}}_t$ and $\tilde{\mathbf{V}}^t$ based on the best sub-optimal solution at step $t-1$. We call this Sequential LPCA (SLPCA):

$$\tilde{\mathbf{a}}_t = \arg \min_{\mathbf{a} \in \mathbb{R}^{1 \times r}} h_t(\mathbf{a}, \tilde{\mathbf{V}}^{t-1}) + \gamma \frac{\|\mathbf{a}\|_F^2}{2}, \forall t \quad (11)$$

$$\tilde{\mathbf{V}}^t = \arg \min_{\mathbf{V} \in \mathbb{R}^{P \times r}} \sum_{s=1}^t h_s(\tilde{\mathbf{a}}_s, \mathbf{V}) + \lambda \frac{\|\mathbf{V}\|_F^2}{2} \quad (12)$$

Equation (11) is easy to solve with a Newton method based gradient descent algorithm, since it is only a vector and the target function is strictly convex. Equation (12) can be solved sequentially based on the past value $\tilde{\mathbf{V}}^{t-1}$. To see how this works, we define a surrogate function $\tilde{h}_t(\mathbf{a}_t, \mathbf{V})$ to approximate $h_t(\mathbf{a}_t, \mathbf{V})$:

$$\begin{aligned} \tilde{h}_t(\tilde{\mathbf{a}}_t, \mathbf{V}) &= h_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^{t-1}) + \nabla_{\mathbf{V}} h_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^{t-1})^T (\mathbf{V} - \tilde{\mathbf{V}}^{t-1}) \\ &\quad + \frac{\alpha_t}{2} \|\mathbf{V} - \tilde{\mathbf{V}}^{t-1}\|_F^2, \quad \alpha_t \geq \|\nabla_{\mathbf{V}}^2 h_t\|_{\text{opt}} \end{aligned} \quad (13)$$

where $\|\cdot\|_{\text{opt}}$ is the operator norm. From the above it follows that $\tilde{h}_t(\tilde{\mathbf{a}}_t, \mathbf{V}) \geq h_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^{t-1})$, and moreover, as we solve Equation (12) under \tilde{h}_t instead of h_t , we get:

$$\tilde{\mathbf{V}}^t = \tilde{\mathbf{V}}^{t-1} - \eta_t \nabla_{\mathbf{V}} h_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^{t-1}) \quad (14)$$

where $\eta_t \propto (\sum_{\tau=1}^t \alpha_\tau)^{-1}$ is the step size. The choice of step size η_t deserves some discussions. We will investigate in Section III on the convergence of this algorithm w.r.t. the BLPCA result. The full SLPCA algorithms is shown below.

```

begin
  Input: data  $\mathbf{X} \in \mathbb{R}^{N \times P}$ ,  $\mathbf{X}^* = 2\mathbf{X} - 1 \in \{-1, 1\}$ ;
  Initialization:  $\tilde{\mathbf{V}}^t \approx 0$ ,  $C, \gamma, \epsilon, \beta \in (0, 1)$ ,  $\alpha$ ;
  for  $t = 1, \dots, N$ ,  $l_t(\tilde{\mathbf{a}}_t) \doteq h_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^{t-1}) + \lambda \frac{\|\tilde{\mathbf{a}}_t\|_F^2}{2}$  do
    Initialize  $\tilde{\mathbf{a}}_t = 0$ ;
    Initialize  $\Delta = \nabla l_t(\tilde{\mathbf{a}}_t) (\nabla^2 l_t(\tilde{\mathbf{a}}_t))^{-1} \nabla l_t(\tilde{\mathbf{a}}_t)$ ;
    while  $\lambda > \epsilon$  do
      Let  $\Delta = -(\nabla^2 l_t(\tilde{\mathbf{a}}_t))^{-1} \nabla l_t(\tilde{\mathbf{a}}_t)$ ,  $d = d_0$ ;
      while
         $\nabla l_t(\tilde{\mathbf{a}}_t + d\Delta) > \nabla l_t(\tilde{\mathbf{a}}_t) + \alpha d \nabla l_t(\tilde{\mathbf{a}}_t)^T \Delta$  do
          Update  $d = \beta d$ ;
        end
      Update  $\tilde{\mathbf{a}}_t = \tilde{\mathbf{a}}_t + d\Delta$ ;
      Update  $\Delta = \nabla l_t(\tilde{\mathbf{a}}_t) (\nabla^2 l_t(\tilde{\mathbf{a}}_t))^{-1} \nabla l_t(\tilde{\mathbf{a}}_t)$ ;
    end
    Set  $\eta_t$ ;
    Update  $\tilde{\mathbf{V}}^t = \tilde{\mathbf{V}}^{t-1} - \eta_t \nabla_{\mathbf{V}} h_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^{t-1})$ 
  end
end

```

Algorithm 1: Sequential LPCA (SLPCA) Pseudo-Code

III. CONVERGENCE ANALYSIS

In this section, we will discuss the convergence property of the SLPCA algorithm. Since our focus is mainly on developing this algorithm for binary data, we will keep our analysis on the Bernoulli random variable, and the loss function is:

$$h_t(\mathbf{a}_t, \mathbf{V}) = \sum_j \log(1 + \exp(-x_{tj}^* \mathbf{a}_t \mathbf{v}_j^T)) \quad (15)$$

where \mathbf{v}_j is the j^{th} row of \mathbf{V} . It is worthy noted that the similar algorithm can be developed in other exponential family random variables.

A. Evaluation Functions

To evaluate BLPCA and SLPCA, we define three important functions that we want to study.

$$\begin{aligned}
C_N(\mathbf{V}^*) &= \frac{1}{N} \sum_{t=1}^N h_t(\mathbf{a}_t^*, \mathbf{V}^*) \\
\hat{C}_N(\tilde{\mathbf{V}}^N) &= \frac{1}{N} \sum_{t=1}^N h_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^N) \\
\tilde{C}_N(\tilde{\mathbf{V}}^N) &= \frac{1}{N} \sum_{t=1}^N \tilde{h}_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^N)
\end{aligned} \quad (16)$$

where $C_N(\mathbf{V}^*)$ is the average batch loss function; $\hat{C}_N(\tilde{\mathbf{V}}^N)$ is the average sequential loss function; $\tilde{C}_N(\tilde{\mathbf{V}}^N)$ is the average sequential loss function under the surrogate function. In [13] we have the relationship:

$$C_N(\mathbf{V}^*) \leq \hat{C}_N(\tilde{\mathbf{V}}^N) \leq \tilde{C}_N(\tilde{\mathbf{V}}^N) \quad (17)$$

In online learning, *Regret* is also of interest [8] [9]. *Regret* takes locally best solution in every step, defined as:

$$\widehat{Re}_N = \frac{1}{N} \sum_{t=1}^N h_t(\tilde{\mathbf{a}}_t, \tilde{\mathbf{V}}^t) \quad (18)$$

B. Convergence Analysis

Lemma 1. For $t = 1, \dots, N$ and $h_t(\cdot)$ defined in (15), $\|\nabla_{\mathbf{V}} h_t\|_F \leq \|\mathbf{a}\|_F$, and $\|\nabla_{\mathbf{V}}^2 h_t\|_{opt} \leq \frac{1}{4} \|\mathbf{a}\|_F^2$.

Proof: For $h_t(\mathbf{a}_t, \mathbf{V})$, w.l.o.g., let $\text{rank}(\Theta) = 1$, we have:

$$\begin{aligned}
[\nabla_{\mathbf{V}} h_t]_j &= -\frac{x_{tj}^* \mathbf{a}_t}{1 + \exp(x_{tj}^* \mathbf{a}_t \mathbf{v}_j^T)} \\
[\nabla_{\mathbf{V}}^2 h_t]_{ij} &= \left(\frac{x_{tj}^* \mathbf{a}_t \delta_{ij}}{2 \cosh(\frac{1}{2} x_{tj}^* \mathbf{a}_t \mathbf{v}_j^T)} \right)^2
\end{aligned}$$

where $\delta_{ij} = 1$ only when $i = j$ means matrix $\nabla_{\mathbf{V}}^2 h_t$ is diagonal. Since $\cosh(x) \geq 1$, hence the norms satisfy $\|\nabla_{\mathbf{V}} h_t\|_F \leq \|\mathbf{a}\|_F$, and $\|\nabla_{\mathbf{V}}^2 h_t\|_{opt} \leq \frac{1}{4} \|\mathbf{a}\|_F^2$. ■

Lemma 2. Let $\tilde{\mathbf{a}}_t$ be bounded by Ω , for $\forall t = 1, \dots, N$. Based on (14) we have $\|\tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1}\|_F \leq \eta_t \Omega$.

Proof: From Equation (14), we have $\|\tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1}\|_F = \eta_t \|\nabla_{\mathbf{V}} h_t\|_F$. Since $\tilde{\mathbf{a}}_t$ result from a regularized problem in (11), so $\tilde{\mathbf{a}}_t$ is bounded by Ω . Thus we have $\|\tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1}\|_F \leq \eta_t \|\tilde{\mathbf{a}}_t\|_F \leq \eta_t \Omega$. ■

Lemma 3. For $h_t(\cdot)$ in Equation (15). $\langle \mathbf{a}, \nabla_{\mathbf{a}} h_t \rangle = \langle \mathbf{V}, \nabla_{\mathbf{V}} h_t \rangle$. Hence, for $t = 1, \dots, N$, $\eta_t \gamma \|\tilde{\mathbf{a}}_t\|_F^2 = \langle \tilde{\mathbf{V}}^{t-1}, -\eta_t \nabla_{\mathbf{V}} h_t \rangle = \langle \tilde{\mathbf{V}}^{t-1}, \tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1} \rangle$.

This follows directly from (11) and (14).

Lemma 4. $h_t(\cdot)$ and surrogate function $\tilde{h}_t(\cdot)$, as well as their first derivative $\nabla h_t(\cdot)$ and $\nabla \tilde{h}_t(\cdot)$ are all Lipschitz continuous.

This is indicated directly from Lemma 1 & Lemma 2 and the definition of Lipschitz continuous [14].

Theorem 1 (Proposition 2, [13]). Under the regularity condition of Lemma 4, and $h_t(\cdot)$ a convex function, $\hat{C}_N(\tilde{\mathbf{V}}^N)$ converges a.s. to $C_N(\mathbf{V}^*)$. Thus, from (17) we directly see that $\tilde{C}_N(\tilde{\mathbf{V}}^N)$ converges a.s. to $C_N(\mathbf{V}^*)$.

The Proof can be found in [15] and [13], following a quasi-martingale theory.

Theorem 2. Given step size as $\eta_t = Ct^{-1/2}$ or $\eta_t = C$, the Regret \widehat{Re}_N converges to within a constant of $\hat{C}_N(\tilde{\mathbf{V}}^N)$, and thus converges to within a constant of $C_N(\mathbf{V}^*)$.

A sketch of proof is given in Appendix A. The results basically show that $\lim_{N \rightarrow \infty} |\widehat{Re}_N - \hat{C}_N(\tilde{\mathbf{V}}^N)| \leq \frac{\gamma \Omega^2}{2}$ if $\eta_t = Ct^{-1/2}$ and $\lim_{N \rightarrow \infty} |\widehat{Re}_N - \hat{C}_N(\tilde{\mathbf{V}}^N)| \leq \gamma \Omega^2 + C \Omega^2$ if $\eta_t = C$, Ω as a constant. From Theorem 1 & Theorem 2, we recognize that both the average sequential function and Regret function converge to within a constant from the average batch optimum.

However, it should be noted here that a better convergence result could be possible, probably by re-design the algorithms, which is one of our future tasks.

IV. EXPERIMENTAL RESULTS

A. Simulated Binary-State System

Firstly, we use simulated binary data to test the performance of our SLPCA algorithm in binary-state system. The generation of correlated Bernoulli sequences is illustrated in [16]. In this work, we focus on the case where $\text{rank}(\Theta) = 1$ since this usually demonstrates the best dimension reduction capability. It should be noted here that the extension to multiple Principal Components is straight-forward following the iterative updating rules in [5].

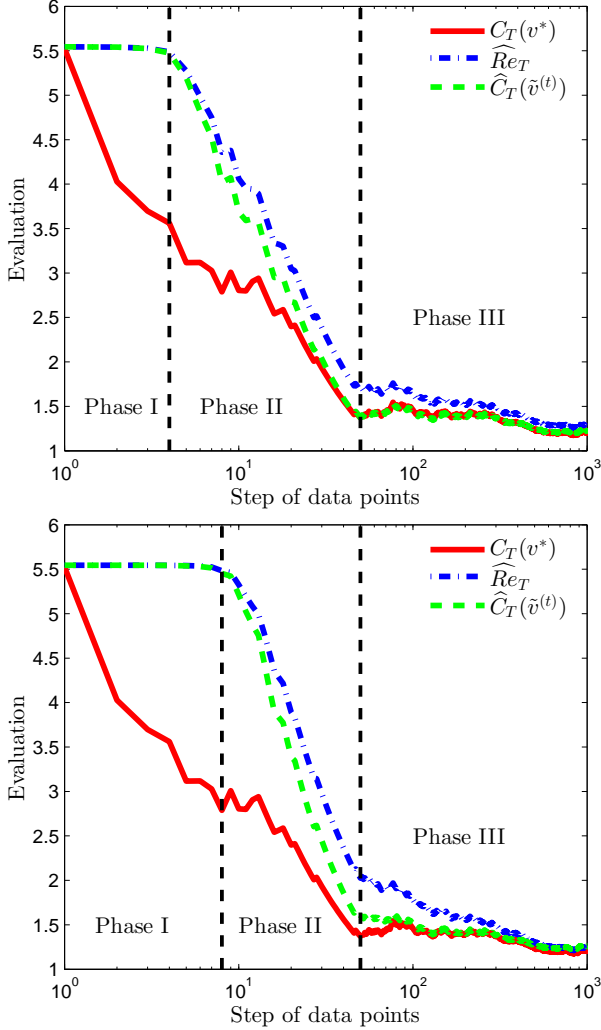


Fig. 1: The three functions $C_t(\mathbf{V}^*)$, $\hat{C}_t(\tilde{\mathbf{V}}^t)$ and \hat{R}_{e_t} as function of t . Top: $\eta_t = Ct^{-1/2}$, with $C = 0.2$, $\gamma = 0.1$. Bottom: $\eta_t = C$, with $C = 0.05$, $\gamma = 0.1$.

We tried the above on data with $P = 8$ dimension and length of $N = 1000$ data points. We initialize $\tilde{\mathbf{V}}^0$ such that its norm is close but not equal to zero, for computation and convergence purposes. Fig 1 shows the three functions defined in (16); whereas Fig 2 shows the key parameters in the sequential steps. There are some interesting findings.

Firstly, though both $\hat{C}_N(\tilde{\mathbf{V}}^N)$ and \hat{R}_{e_N} converges at least within a constant to $C_N(\mathbf{V}^*)$, the stochastic learning can be clearly divided into three Phases, as shown in Fig 1. Phase

I stands for the period when the norm of $\tilde{\mathbf{V}}^0$ is close to zero right after the initialization, when $h_t(\mathbf{a}_t, \mathbf{V})$ approaches $P \log 2$ as in Equation (15). Phase II characterizes the decay of error versus N , whereas Phase III stands for when the error converges to within a constant independent of N .

Secondly, $\|\tilde{\mathbf{V}}^t\|_F^2$ increases versus t , which means that $\|\tilde{\mathbf{V}}^t\|_F^2$ behaves differently from the coefficient in sequential learning of linear model [15] [13]. Matrix factorization places no constraints for $\tilde{\mathbf{V}}^t$, hence cannot guarantee the bound of $\tilde{\mathbf{V}}^t$. From another perspective, $\tilde{\mathbf{a}}_t$ is bounded since Equation (11) has consistent regularization, while $\tilde{\mathbf{V}}_t$ not since there is a summation of loss functions in Equation (12). It should be noted that, in Fig 2, $\tilde{\mathbf{a}}_t$ decreases versus t , which could result from (11) and is an interesting topic in the future.

Thirdly, due to the unbounded $\tilde{\mathbf{V}}^t$, the term $\|\tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1}\|_F$ is not $\propto t^{-1}$ as in [15] and [13]. It should be noted that the theoretical bound for $\|\tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1}\|_F$ under constant step size could be as low as $t^{-1/2}$, which could be a result of the convergence behavior of $\tilde{\mathbf{a}}_t$ under constant step size.

Last but not least, it is important to mention that the bounds obtained in Theorem 2 assume N large enough. However, in many cases the decay of N is not that fast. Therefore, the effect of N cannot be completely ignored in the analysis.

B. Building End-Use Energy Modeling

Here, we introduce an application of SLPCA in Building Energy End-Use Modeling. Building End-Uses corresponds to the energy sectors that are occupant-driven. This subject has attracted significant interest in recent years because building energy shows strong dependence on end-user behavior, e.g. plug-in loads, user-controlled lighting, user-adjusted HVAC, etc. [17] [18].

Energy end-use modeling has been attempted from either a top-down or a bottom-up approach. In this work, since we are more interested in modeling occupant behavior, we adopt the bottom-up approach. This approach is usually based on stochastic simulations of the energy usage pattern for each individual appliance. Dimensional reduction can help to generate one or more *Principal Appliances*, and can more efficiently characterize the whole space energy consumption.

Here, we want to study the modeling of all the computer monitors in a small, shared work space. We collect the data of 6 monitors in 10 minutes interval, and use BLPCA and SLPCA to obtain the *Principal Monitor* profile of the building. Considering that the pattern could be non-stationary, we choose the constant step size that is short enough to track the changes as they appear¹. We also only consider the first *Principal Monitor* to achieve the best dimensional reduction. The convergence of the algorithm is shown in Fig 3. We observe a good convergence for both $\hat{C}_N(\tilde{\mathbf{V}}^N)$ and \hat{R}_{e_N} . Periodic fluctuation is observed, due to the periodic transition between day and night energy consumption, which results in periodical changing of the data model. Moreover, the online algorithm demonstrate less fluctuations because they adaptively update the model of the data.

¹one could presumably also leverage the likely periodic behavior of the data by appropriate aggregation

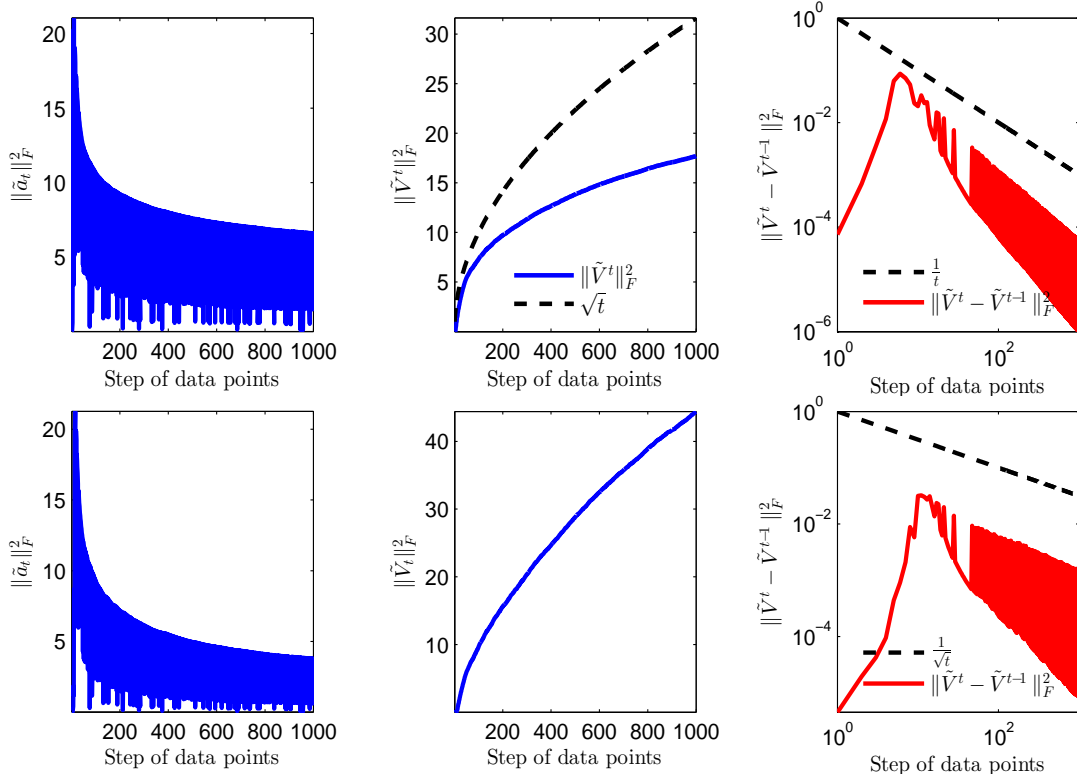


Fig. 2: The convergence property of $\tilde{\mathbf{a}}_t$, $\tilde{\mathbf{V}}^t$ and $\|\tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1}\|_F$. Top: $\eta_t = Ct^{-1/2}$, with $C = 0.2$, $\gamma = 0.1$. Bottom: $\eta_t = C$, with $C = 0.05$, $\gamma = 0.1$.




Fig. 3: The three functions $C_t(\mathbf{V}^*)$, $\widehat{C}_t(\tilde{\mathbf{V}}^t)$ and \widehat{Re}_t as function of t for energy end-use simulation with constant step size $\eta_t = C$ as $C = 0.05$, $\gamma = 0.1$.

The BLPCA, SLPCA and Regret are used to reconstruct the original data, as illustrated in Section II-C, when we discussed the reconstruction of \mathbf{X} by $g(\Theta)$. The results are compared with the original data in Fig 4 (sum of states of all appliances, 1 as ON and 0 as OFF). Interestingly, Regret gives better approximation to BLPCA since it uses locally best pairs of $\tilde{\mathbf{a}}_t$ and $\tilde{\mathbf{V}}^t$, so that can better catch the periodic pattern of the original data. Whereas SLPCA uses the $\tilde{\mathbf{V}}^T$, which could probably give unpromising result if data is non-stationary.

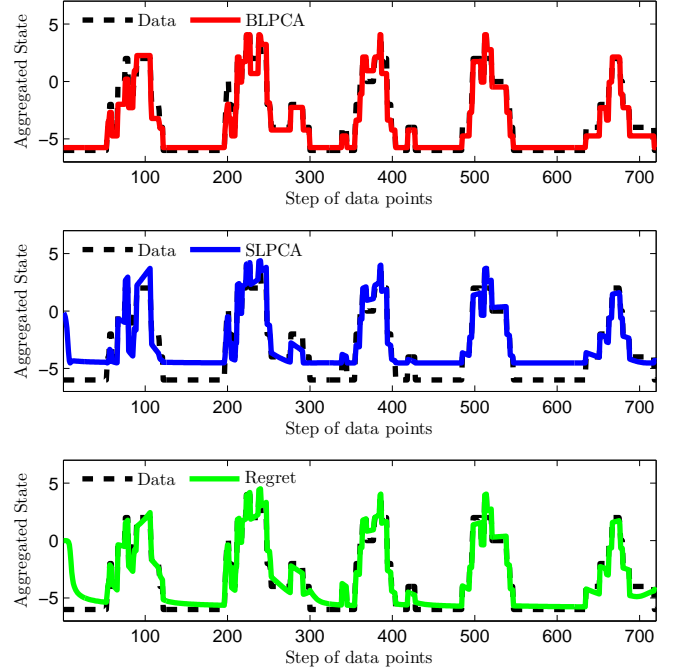


Fig. 4: Reconstruction of the aggregated state (sum of states of 6 monitors) under the three functions $C_t(\mathbf{V}^*)$, $\widehat{C}_t(\tilde{\mathbf{V}}^t)$ and \widehat{Re}_t as function of t .

V. CONCLUSION

Sequential or online dimension reduction addresses more and more attentions due to the explosion of streaming data based application and the requirement of adaptive statistical

modeling in many emerging fields. In this work, we extend the theory of ePCA or LPCA to sequential version based on online convex optimization theory, which can maintain the capability to model large families of distributions, at the same time achieve the computation and storage efficiency. In our work, we define two functions to evaluate the SLPCA algorithm, the average sequential target function $\widehat{C}_N(\tilde{\mathbf{V}}^N)$ and the Regret function \widehat{Re}_N , and show that both of them converge at least within a constant to BLPCA results. We also demonstrate an application of this algorithm in building energy end-use modeling.

APPENDIX

Lemma 5. For $t = 1, \dots, N$, if Ω is the upper bound of $\|\mathbf{a}\|_{opt}^2$ as in Lemma 2, $\|\tilde{\mathbf{V}}^t\|_F^2 \leq \Omega^2 \sum_{s=1}^t \eta_s^2 + 2\gamma\Omega^2 \sum_{s=1}^t \eta_s$.

Proof: We start from the relationship:

$$\begin{aligned} \|\tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1}\|_F^2 &= \|\tilde{\mathbf{V}}^t\|_F^2 - \|\tilde{\mathbf{V}}^{t-1}\|_F^2 - 2\langle \tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^{t-1}, \tilde{\mathbf{V}}^{t-1} \rangle \\ &= \|\tilde{\mathbf{V}}^t\|_F^2 - \|\tilde{\mathbf{V}}^{t-1}\|_F^2 - 2\eta_t \gamma \|\tilde{\mathbf{a}}^t\|_F^2 \end{aligned}$$

We sum over the LHS and RHS and get:

$$\sum_{s=1}^t \|\tilde{\mathbf{V}}^s - \tilde{\mathbf{V}}^{s-1}\|_F^2 + 2\gamma \sum_{s=1}^t \eta_s \|\tilde{\mathbf{a}}^s\|_F^2 = \|\tilde{\mathbf{V}}^t\|_F^2 - \|\tilde{\mathbf{V}}^0\|_F^2$$

For simplicity, assume $\|\tilde{\mathbf{V}}^0\|_F^2 \approx 0$, we prove the lemma. ■

Now turn to proof of Theorem 2. Based on (14) we have:

$$\begin{aligned} \|\tilde{\mathbf{V}}^t - \tilde{\mathbf{V}}^N\|_F^2 &= \|\tilde{\mathbf{V}}^{t-1} - \tilde{\mathbf{V}}^N\|_F^2 + \eta_t^2 \|\nabla_{\mathbf{V}} h_t\|_F^2 \\ &\quad - 2\eta_t \langle \nabla_{\mathbf{V}} h_t, \tilde{\mathbf{V}}^{t-1} - \tilde{\mathbf{V}}^N \rangle \end{aligned}$$

From Lemma 1, Lemma 5, and $\|\nabla_{\mathbf{V}} h_t\|_F^2 \leq \Omega^2$, thus:

$$\begin{aligned} N\{\widehat{Re}_N - \widehat{C}_N(\tilde{\mathbf{V}}^N)\} &\leq \sum_{t=1}^N \langle \nabla_{\mathbf{V}} h_t, \tilde{\mathbf{V}}^{t-1} - \tilde{\mathbf{V}}^N \rangle \\ &\leq \frac{\|\tilde{\mathbf{V}}^N\|_F^2}{2\eta_0} + \sum_{t=1}^N \left(\frac{1}{2\eta_t} - \frac{1}{2\eta_{t-1}} \right) \|\tilde{\mathbf{V}}^N - \tilde{\mathbf{V}}^{t-1}\|_F^2 + \frac{\Omega^2}{2} \eta_t \\ &\leq \frac{\|\tilde{\mathbf{V}}^N\|_F^2}{2\eta_0} + \sum_{t=1}^N \left(\frac{1}{2\eta_t} - \frac{1}{2\eta_{t-1}} \right) \|\tilde{\mathbf{V}}^N\|_F^2 + \frac{\Omega^2}{2} \eta_t \end{aligned}$$

- diminishing step size $\eta_t = Ct^{-1/2}$. From Lemma 5, we have:

$$\begin{aligned} |\widehat{Re}_N - \widehat{C}_N(\tilde{\mathbf{V}}^N)| &\leq \frac{\Omega^2 C \log N}{2} \frac{1}{N} + \frac{\Omega^2 C \log N}{4} \frac{1}{\sqrt{N}} \\ &\quad + \frac{\Omega^2(2\gamma + C)}{2\sqrt{N}} + \frac{\gamma\Omega^2}{2} \end{aligned}$$

Then $\lim_{N \rightarrow \infty} |\widehat{Re}_N - \widehat{C}_N(\tilde{\mathbf{V}}^N)| \leq \frac{\gamma\Omega^2}{2}$. But with reasonable N , the term $\frac{\Omega^2 C \log N}{\sqrt{N}}$ will also be significant. Usually, small C and γ can force a lower error bound. However, small γ can result in more steps in optimizing for $\tilde{\mathbf{a}}_t$, whereas small C would make the step size too small, which may not be a good choice if we want a fast decaying of the error bound.

- constant step size $\eta_t = C$: For constant step, we have:

$$|\widehat{Re}_N - \widehat{C}_N(\tilde{\mathbf{V}}^N)| \leq \gamma\Omega^2 + \Omega^2 C$$

Similarly, we prefer small C and γ . The challenge of using small C and γ have already been discussed.

ACKNOWLEDGMENT

This research is funded by the Republic of Singapore National Research Foundation through a grant to the Berkeley Education Alliance for Research in Singapore (BEARS) for the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) Program. BEARS has been established by the University of California, Berkeley as a center for intellectual excellence in research and education in Singapore.

REFERENCES

- [1] A. I. Schein, L. K. Saul, and L. H. Ungar, "A generalized linear model for principal component analysis of binary data," in *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, vol. 38, 2003, p. 46.
- [2] T. Hastie, R. Tibshirani, J. Friedman, T. Hastie, J. Friedman, and R. Tibshirani, *The elements of statistical learning*. Springer, 2009, vol. 2, no. 1.
- [3] S. Papadimitriou, J. Sun, and C. Faloutsos, "Streaming pattern discovery in multiple time-series," in *Proceedings of the 31st international conference on Very large data bases*. VLDB Endowment, 2005, pp. 697–708.
- [4] M. E. Tipping and C. M. Bishop, "Probabilistic principal component analysis," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 61, no. 3, pp. 611–622, 1999.
- [5] M. Collins, S. Dasgupta, and R. E. Schapire, "A generalization of principal components analysis to the exponential family," in *Advances in neural information processing systems*, 2001, pp. 617–624.
- [6] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with bregman divergences," *The Journal of Machine Learning Research*, vol. 6, pp. 1705–1749, 2005.
- [7] L. M. Bregman, "The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming," *USSR computational mathematics and mathematical physics*, vol. 7, no. 3, pp. 200–217, 1967.
- [8] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," 2003.
- [9] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2011.
- [10] R. Vidal, Y. Ma, and S. Sastry, "Generalized principal component analysis (gpca)," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 27, no. 12, pp. 1945–1959, 2005.
- [11] G. Casella and R. L. Berger, *Statistical inference*. Duxbury Press Belmont, CA, 1990, vol. 70.
- [12] I. Csisz, G. Tusnady *et al.*, "Information geometry and alternating minimization procedures," *Statistics and decisions*, 1984.
- [13] M. Mardani, G. Mateos, and G. B. Giannakis, "Rank minimization for subspace tracking from incomplete data," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 5681–5685.
- [14] D. P. Bertsekas, "Nonlinear programming," 1999.
- [15] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *The Journal of Machine Learning Research*, vol. 11, pp. 19–60, 2010.
- [16] A. D. Lunn and S. J. Davies, "A note on generating correlated binary variables," *Biometrika*, vol. 85, no. 2, pp. 487–490, 1998.
- [17] L. G. Swan and V. I. Ugursal, "Modeling of end-use energy consumption in the residential sector: A review of modeling techniques," *Renewable and Sustainable Energy Reviews*, vol. 13, no. 8, pp. 1819–1835, 2009.
- [18] Z. Kang, M. Jin, and C. J. Spanos, "Modeling of end-use energy profile: An appliance-data-driven stochastic approach," *arXiv preprint arXiv:1406.6133*, 2014.